Re: AI Accountability Policy Request for Comment (230407-0093)

Lumeris appreciates the opportunity to respond to the National Telecommunications and Information Administration's (NTIA) *AI Accountability Policy Request for Comment*. Lumeris is a leading provider of technology, insurance capabilities and expertise to support value-based care models in the healthcare industry. We are dedicated to the idea that radical change in health outcomes and performance occurs by placing patients at the center of care and decisioning, support by committed provider, payers, and technologies moving synergistically toward the goal of improving and maintaining a patient's health. Our commitment to our health plan affiliate Essence members is the same – high quality care can be accessible and affordable to all beneficiaries when supported by value-based agreements that drive both quality outcome measures and cost management.

As a health care solutions company, we are dedicated to deploying and leveraging technology solutions to improve health outcomes and optimize provider practices. Our technology suite – Lumeris**Realize** (data transformation, analytics, and insights), Lumeris**Engage** (care orchestration and patient engagement), and Lumeris**Protect** (health plan operations and management) – work in concert to enable better clinical decision making through data, improve patient outcomes through meaningful AI-driven engagement and health co-management between patient and provider, and drive the highest quality outcomes for consumers through health plan performance. As an industry leader in and champion of the value-based care initiative, our primary focus resides in developing tools to drive health outcomes, leveraging over a decade of experience as a collaborative payer through our operation of Essence Healthcare, a payer first envisioned by providers.

We applaud the NTIA for their dedication to ensuring efficacious and secure usage of AI technologies as a solution for a variety of challenges faced by various sectors, including the health care industry. We welcome the agency's openness to stakeholder input and ongoing commitment to legal, effective, safe, and trustworthy use of technologies. Below is feedback from Lumeris team members and key stakeholders regarding various provisions of the rule. The original questions from the request for comment are indicated by italicized and/or bolded text.

AI Accountability Objectives

What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments?

> Lumeris recommends that AI accountability mechanisms should focus on the systems and decisions which are driven by the AI. Additionally, assessments provide a powerful tool for ensuring accuracy of AI tools and compliance with applicable regulations and ethical considerations. Internal assessments are most valuable for communicating the importance of ensuring that systems are not biased. External assessments are most valuable for challenging the existing norms and processes in use at the company.

AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there tradeoffs among these goals? To what extent can these inquiries be conducted by a single team or instrument?

> ***The AI system does not substantially contribute to harmful discrimination against people.*** This is fundamentally the most essential goal of any accountability measure. It is of paramount importance that AI systems do not result in discrimination, either implicitly or explicitly. There are fundamental statistical tests which can be performed to demonstrate whether a given model is exhibiting harmful discrimination. Even this type of analysis requires additional review because many diseases have a

genetic component which may be correlated to race/ethnicity. For instance, if an AI system is meant to identify patients likely to have symptoms of sickle-cell disease, it would be appropriate for that system to disproportionately identify African American patients compared to other demographic patients due to the genetic variation in the disease distribution.

The AI system does not substantially contribute to harmful misinformation, disinformation, and other forms of distortion and content-related harms. This requirement has the possibility to be fundamentally impossible and a limitation placed uniquely on AI systems in the healthcare industry. Clinical guidance often includes old, out-of-date, and even erroneous information and guidance that does not rise to the level of malpractice. A real-world example is the over-prescription of antibiotics. By the CDC's estimates, nearly 50% of antibiotics are prescribed inappropriately. Despite poor performance, there has not been a serious dialogue about misinformation and disinformation that clinicians contribute to.

It is a reasonable concern that requiring AI systems to not contribute to harmful misinformation would require that AI systems attain a level of perfection and accuracy that exceed our current reliance on clinical judgement. This is counterproductive as AI systems are easier to revise and correct as clinical practices evolve. A well-designed and implemented AI system which managed access to antibiotics, could significantly reduce adverse outcomes and the spread of antibiotic-resistant "super bugs." Such systems have the potential to produce more accurate, though still imperfect, guidance about the use of antibiotics.

The AI system protects privacy. Within the Healthcare Industry, this requirement is addressed by the Health Insurance Portability and Accountability Act (HIPAA) and other existing regulations that place control over the access to personal data and give individuals the opportunity to manage access to and use of their healthcare data. At least within this area, any additional requirements would risk causing confusion about how to apply multiple and overlapping regulations.

There has been adequate transparency and explanation to affected people about the uses, capabilities, and limitations of the AI system. Transparency and explanation are valuable; however, the value of AI systems can be very specific and hard to explain.

There are adequate human alternatives, consideration, and fallbacks in place throughout the AI system lifecycle. Legislation requiring human alternatives to AI systems will result in worse outcomes for patients. Take, for instance, the recent paper about the NYUTron AI model which predicts which patients are at greatest risk of being readmitted to the hospital. In this paper, the authors specifically compared the performance of physicians and the AI model at predicting which patients would be readmitted and found that the model was more accurate. In this environment, any human judgement (even clinical expertise) which is included in the system will result in fewer readmissions being targeted and some patients never likely to readmit would be included in a program. Combined, these effects will worsen outcomes and cost more money.

There has been adequate consultation with, and there are adequate means of contestation and redress for, individuals affected by AI system outputs. The focus on managing access to AI systems necessarily must be placed on balancing the need of providing access with implementing safeguards that limit access. For instance, we truly believe that AI systems present a significant opportunity to improve the delivery and outcomes of healthcare. As such, it is important to ensure that patients always

have the option to have access to AI systems to provide new sources of data to the AI system to improve the system's ability to improve their care. As these systems improve, denying patients access to the systems and their results will look more and more like denying access to a scalpel or medication.

There is adequate management within the entity deploying the AI system such that there are clear lines of responsibility and appropriate skillsets. This kind of validation of the process for building AI systems sounds reasonable. For instance, one could ask "was a clinician consulted when an AI system to reduce inpatient admissions was developed?"

The application of accountability measures (whether voluntary or regulatory) is more straightforward for some trustworthy AI goals than for others. With respect to which trustworthy AI goals are there existing requirements or standards? Are there any trustworthy AI goals that are not amenable to requirements or standards? How should accountability policies, whether governmental or non-governmental, treat these differences?

Within the healthcare environment, there are several existing requirements and standards which provide some form of guidance for trustworthy AI.

Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?

The concept of "explainability" presents fundamental challenges. Within healthcare, many AI models can demonstrate their value at predicting risk of complications or other outcomes. However, explaining why those models are effective is extremely challenging and explanations are typically developed retrospectively after the fact so that even a definition of what meets the criteria of "explain" becomes almost impossible.

Existing Resources and Models

What AI accountability mechanisms are currently being used? Are the accountability frameworks of certain sectors, industries, or market participants especially mature as compared to others? Which industry, civil society, or governmental accountability instruments, guidelines, or policies are most appropriate for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?

Within healthcare, it is not uncommon to evaluate AI systems to ensure they do not unintentionally encode a gender, racial, or ethnic bias. This requires a review of results as the system is prepared and continuing evaluation of the system's performance. Such work is typically conducted by the teams that create those systems.

AI Accountability Policies

What role should government policy have, if any, in the AI accountability ecosystem?

Should AI accountability policies and/or regulation be sectoral or horizontal, or some combination of the two? A number of the challenges and risks are distinct to the individual sector of the solution, which the accountability policies must uniquely reflect. We are most familiar with the healthcare sector which already includes a number of limitations on the availability of member information, the uses of that information, and the member's ability to audit and limit the uses of their information.

Should AI accountability regulation, if any, focus on inputs to audits or assessments (e.g., documentation, data management, testing and validation), on increasing access to AI systems for auditors and researchers, on mandating accountability measures, and/or on some other aspect of the accountability ecosystem? We believe it is more important to focus on inputs to audits and assessments rather than on increasing access to auditors. Ultimately, within a private organization, it should be the responsibility of the board of directors (through their compliance team) to ensure audits or assessments are conducted as expected and according to generally acceptable practices.

How can government work with the private sector to incentivize the best documentation practices?

Clarity about expectations is absolutely vital. Without a common understanding of what "the best documentation practices" are it is unlikely that that actors will have a consistent set of practices and documentation.

Is it important that there be uniformity of AI accountability requirements and/or practices across the United States? Across global jurisdictions? If so, is it important only within a sector or across sectors? What is the best way to achieve it? Alternatively, is harmonization or interoperability sufficient and what is the best way to achieve that?

As a result of the different existing regulations for various sectors of the business community and the differential impact on individual's lives, it makes sense for the regulations to occur primarily within sectors rather than broadly across sectors. For instance, the impact of AI systems for managing retirement funds or systems for managing healthcare treatments is much more significant than the impact that AI systems which conduct transactions at drive-thru windows. Both the severity of potential outcomes and risks are so fundamentally different that a single accountability requirement would be of limited value. To the extent there are published notifications about the results of accountability efforts, these need to be focused on those areas where those models have the most significant potential impact on an individual. The more ubiquitous these notifications are the more likely their messages are to be limited or watered down. However, we do feel that its important to address AI requirements for various sectors in a uniform manner across the United States to avoid conflicting state laws. One way to achieve uniform requirements across the United States is to finalize federal regulation that specifically pre-empts any state law on the subject matter, similar to the manner in which the federal Medicare Advantage statutes preempt most state laws that would otherwise govern those plans.